

APPLICATION FOR UNITED STATES LETTERS PATENT

For

PRE-SUMMARIZATION AND ANALYSIS OF RESULTS GENERATED BY AN AGENT

Inventors:

Vijay Deshmukh

Benjamin Swartzlander

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP

32400 Wilshire Boulevard

Los Angeles, CA 90025-1026

(408) 720-8300

Attorney's Docket No.: 005693.P043

"Express Mail" mailing label number: EV409365110US

Date of Deposit: March 12, 2004

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria VA 22313-1450

Patricia Richard
(Typed or printed name of person mailing paper or fee)

[Signature]
(Signature of person mailing paper or fee)

3/12/2004
(Date signed)

Pre-summarization and Analysis of Results Generated by an Agent

FIELD OF THE INVENTION

[0001] At least one embodiment of the present invention pertains to networked storage systems, and more particularly to a method and apparatus for collecting and reporting data pertaining to files stored on a storage server.

BACKGROUND

[0002] A file server is a type of storage server which operates on behalf of one or more clients to store and manage shared files in a set of mass storage devices, such as magnetic or optical storage based disks. The mass storage devices are typically organized as one or more groups of Redundant Array of Independent (or Inexpensive) Disks (RAID). One configuration in which file servers can be used is a network attached storage (NAS) configuration. In a NAS configuration, a file server can be implemented in the form of an appliance, called a filer, that attaches to a network, such as a local area network (LAN) or a corporate intranet. An example of such an appliance is any of the NetApp Filer products made by Network Appliance, Inc. in Sunnyvale, California.

[0003] A filer may be connected to a network, and may serve as a storage device for several users, or clients, of the network. For example, the filer may store user directories and files for a corporate or other network, such as a LAN or a wide area network (WAN). Users of the network can be assigned an individual directory in which they can store personal files. A user's directory can then be accessed from computers connected to the network.

[0004] A system administrator can maintain the filer, ensuring that the filer continues to have adequate free space, that certain users are not monopolizing storage on the filer, etc. A Multi-Appliance Management Application (MMA) can be used to monitor the storage on the filer. An example of such an MMA is the Data Fabric Monitor (DFM) products made by Network Appliance, Inc. in Sunnyvale, California. The MMA may provide a Graphical User Interface (GUI) that allows the administrator to more easily observe the condition of the filer.

[0005] The MMA needs to collect information about files stored on the filer to report back to the administrator. This typically involves a scan, also referred to as a “file walk” of storage on the filer. During the file walk, the MMA can determine characteristics of files stored on the filer, as well as a basic structure, or directory tree, of the directories stored thereon. These results can be accumulated, sorted, and stored in a database, where the administrator can later access them. The MMA may also summarize the results of the file walk so they are more easily readable and understood by the administrator.

[0006] On a large system, the file walk can be a very intensive process. Additionally, the results of a typical file walk can themselves be very large and difficult to parse. An MMA typically has many tasks to perform, and generally should be available for the administrator. What is needed is a way to reduce the load on an MMA while still maintaining and monitoring attached appliances.

SUMMARY OF THE INVENTION

[0007] A method for collecting information from a storage server is disclosed.

An agent scans a storage server. Information regarding files stored on the storage server is collected. The agent then summarizes the information, creating a summary. The summary is stored on a database server.

[0008] Other aspects of the invention will be apparent from the accompanying figures and from the detailed description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] One or more embodiments of the present invention are illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0010] **Figure 1** illustrates a monitoring system for a storage server;

[0011] **Figure 2** illustrates a block diagram of an agent;

[0012] **Figure 3** is a flowchart illustrating a process for pre-summarizing and analyzing results generated by an agent;

[0013] **Figure 4** illustrates a table displaying a list of interesting files;

[0014] **Figure 5** illustrates a table listing information about directories on the server;

[0015] **Figure 6** illustrates a histogram showing server usage of certain users; and

[0016] **Figure 7** illustrates a histogram showing the types of files stored on a server.

DETAILED DESCRIPTION

[0017] Described herein are methods and apparatuses for Pre-summarization and Analysis of Results Generated by an Agent. Note that in this description, references to “one embodiment” or “an embodiment” mean that the feature being referred to is included in at least one embodiment of the present invention. Further, separate references to “one embodiment” or “an embodiment” in this description do not necessarily refer to the same embodiment; however, such embodiments are also not mutually exclusive unless so stated, and except as will be readily apparent to those skilled in the art from the description. For example, a feature, structure, act, etc. described in one embodiment may also be included in other embodiments. Thus, the present invention can include a variety of combinations and/or integrations of the embodiments described herein.

[0018] According to an embodiment of the invention, an agent is coupled to a storage server through a network. The storage server is monitored by a Multi-Appliance Management Application (MMA). The agent performs a scan, or a “file walk,” of the storage server and returns the results to the MMA through the network. The results can then be stored on a database server. The agent is responsible for collecting information about files stored on the storage server. The agent is also responsible for generating summaries, including tables and histograms, of relevant and requested information about the files on the server before the information is transferred to the MMA. In this way, the agent pre-summarizes the information before it is transmitted to the MMA. As a result, the MMA is not burdened with the task of summarizing the information, and the summaries are available to an administrator as soon as they are requested.

[0019] The MMA is generally a single server that is used to allow a system administrator to monitor a storage or file server. When a large storage server is monitored, the MMA may have difficulty performing its monitoring duties and a file walk at the same time. In fact, the file walk may make the MMA inaccessible to the system administrator, and the MMA may become a bottleneck, since it may be incapable of performing the file walk in a reasonable amount of time. According to an embodiment of the invention, independent agents are used to perform the file walk, to reduce the load on the MMA. At a later time, the system administrator may want summarized information about the file server. Instead of having the MMA summarize the information, the summaries are compiled by the agent during the file walk, and stored on the database.

[0020] **Figure 1** illustrates a monitoring system for a storage server. The system 100 includes a filer 102, an MMA 104 including a monitor 106, a database 108, a graphical user interface (GUI) 110, and two agents 112 and 114. The agents 112 and 114 can perform a file walk of the filer 102 for the MMA 104. An agent may be an independent server that is attached to the network and is dedicated to performing file walks. By having an agent perform this task rather than having the MMA do it, the MMA can save its resources for other tasks, such as monitoring current activity on the filer 102 using the monitor 106. Ultimately, one goal is to minimize the amount of work the MMA is required to do. Additionally, multiple agents can be added to perform a complete file walk in less time if necessary.

[0021] According to one embodiment of the invention, the agents 112 and 114 may use a file system different from the one used by the filer 102. For example, the agent 112

uses the Common Internet File System (CIFS), while the agent 114 uses the Network File System (NFS). Here, either agent 112 or 114 is able to perform the file walk of the filer 102, regardless of the file system used by the filer 102. The agent 112 also has storage 116 to store the results of a file walk while the walk is occurring and before they are transferred to the MMA 104. The agent 114 may also have attached storage for this purpose.

[0022] The filer 102 is generally attached to a volume 118. The volume 118 may include one or more physical hard drives or removable storage drives that comprise the storage for the filer 102. For example, the volume 118 may comprise a RAID structure. The filer 102 may also be connected to other volumes that comprise storage. A file walk generally scans all files stored on the entire volume 118, regardless of whether all of the files are stored on the same physical drive. Further, although the volume 118 may contain several separate physical drives, the volume 118 may appear and function as a single entity.

[0023] The results of a file walk may be transferred to and stored on the database server 108 after the file walk is complete. The database server 108 can then be accessed by the GUI 110, so that an administrator can search the results of the file walk. The GUI 110 may allow the administrator to easily parse the results of a specific file walk, including allowing the administrator to monitor the total size of files stored on the filer, the size of particular directories and their subdirectories, the parents of specific directories, etc. These queries will be discussed in more detail below. The file walk may also collect statistics about the files on the filer, such as the total size of files, the most

accessed files, the types of files being stored, etc. According to one embodiment, the GUI 110 may be a web-based Java application.

[0024] According to an embodiment of the invention, the summary is written to the database server 108 as a table or a histogram. The summary may then be accessed through a Java applet using a web browser such as Internet Explorer or Netscape. In another embodiment, the summaries are accessed using other programs. Although tables and histograms are shown here, it is understood that any appropriate manner of relaying the summary data to the administrator may be used.

[0025] **Figure 2** illustrates a block diagram of an agent. The agent 112 includes a processor 202, a memory 204, a network adapter 206, and a storage adapter 208. These components are linked through a bus 210. The agent 112, as shown in **Figure 2**, is typical of a network server or appliance, and it is understood that various different configurations may be used in its place. The agent 114 may be similar.

[0026] The processor 202 may be any appropriate microprocessor or central processing unit (CPU), such as those manufactured by Intel or Motorola. The memory 204 may include a main random access memory (RAM), as well as other memories including read only memories (ROM), flash memories, etc. The operating system 212 is stored in the memory 212 while the agent 112 is operating. The operating system includes the file system, and may be any operating system, such as a Unix or Windows based system. The network adapter 206 allows the agent 112 to communicate with remote computers over the network 214. Here, the agent 112 will be collecting data from the filer 102 and sending data to the MMA 104. The storage adapter 208 allows the agent 112 to communicate with the storage 116 and other external storage.

[0027] **Figure 3** is a flowchart illustrating a process for pre summarizing and analyzing results generated by an agent. In block 302, an agent 112 scans a storage server, such as a filer 102. In one embodiment of the invention, many agents may scan different sections of the volume 118. The MMA 104 may determine how to divide the file walking task among the various agents. In one embodiment, for example, the MMA 104 may assign certain directories within the root directory to a first agent, while the other directories are assigned to a second agent. The MMA 104 may use as many agents as necessary to perform the file walk. For example, when scanning a very large volume 118, several agents may be necessary to perform the file walk in an acceptable time. As a further example, the administrator may want to perform the file walk very quickly, and may assign additional agents to expedite the task.

[0028] During the file walk scanning, in block 304, the agent 112 collects information about files stored on the volume 118. This information may include file names, directory names, file sizes, dates of creation, etc. The file walk may be performed by one or more 'threads.' A thread may be a program capable of operating independently of other programs. Using a single threaded system, the agent scans directories found on the volume 118 with a single thread. A multi-threaded system may include two or more threads. A file thread can be used to scan and determine characteristics of files, while a directory thread can be used to determine the contents of directories. A directory queue and a file queue are also established. The directory thread examines the directory found at the top of the directory queue, and places that directory's contents into the file queue. The file thread then examines the members of the file queue, placing directories in the directory queue and examining files. The file thread may collect information including

the name of the file, the size of the file, the location of the file, the type of file, the time of creation of the file, the time of last access of the file, and the owner of the file. This information will be used to create the tables and histograms in **Figures 4-7**. The directory thread may also report information about the directory structure on the volume 118.

[0029] In block 306, the agent summarizes the collected information and creates tables and histograms. Examples of the summarized information will be shown in **Figures 4-7**. There are several types of summaries that the agent can create. For example, the agent can create a table of interesting files, a table of directory or user information, histograms listing the types of files stored, etc. In block 308, the summarized information is stored on the database server 108. The GUI 110 may be used to later access the stored information.

[0030] **Figure 4** illustrates a table displaying a list of interesting files. While collecting the file data, the agent 112 may keep track of certain statistics about the files on the storage server 102. For example, the table 400 includes a list of several files that the agent 112 has tracked. The agent 112 has been instructed to keep track of the largest file found, the smallest file found, the least recently accessed file found, and the oldest file found. Although these types of files are listed, it is understood that any characteristic may be tracked. For example, the agent 112 may also track the most accessed file, the directory with the largest number of files, etc. The table 400 may also include the top 'n' files of each type, where 'n' is a number specified by the administrator or the MMA 104. The table 400 may be configured so that the GUI 110 can access and display its contents.

[0031] Summaries are useful for a number of reasons. The summaries can give an administrator a quick overview of the operation of the filer 102. The summaries can also point out trouble spots or potential trouble spots. An administrator needs a quick and easy way to monitor a filer 102, and the summaries can be tailored to provide important information. Since the volume 118 may be very large, containing hundreds of thousands or millions of files, it may be impractical for the MMA 104 to summarize the file walk metadata. Therefore, the agents 112 and 114 can generate the summaries while the file walk is occurring. It is easy to add more agents if necessary to cope with the additional workload created by the generation of the summaries. By shifting the summarization task to the agents 112 and 114, the MMA 104 will be more responsive to the requests of the administrator.

[0032] A row 402 lists the name and last date of access of the least recently accessed file on the storage server 102. This information is useful if a system administrator is trying to determine whether there are any old or unused files on the volume 118. For example, if the least recently accessed file was accessed less than six months ago, the administrator may determine that no corrective action is necessary. However, as shown here, there is at least one file that has not been accessed for several years. The table 400 may be configured so that several other old files may also be listed. For example, the table may list any file that has not been accessed in the last year. The administrator can then make a determination about whether the file should be purged or retained. The administrator can determine what to do with these old files. For example, the administrator may delete them or move them to another storage server. In one embodiment, these actions are automated.

[0033] The row 404 lists the largest file found on the filer 102. This information may be useful to an administrator who needs to create or maintain free space on the server, and is looking for large files to remove. Here, a user is storing a very large movie file, which is occupying a sizeable percentage of the server's storage space. The administrator can target this file, deleting it if necessary. The administrator can also configure the agent 112 to include a list of several of the largest files found on the volume.

[0034] The row 406 lists the smallest file found. The row 408 lists the oldest file found. This information may be useful to an administrator trying to determine what type of usage occurring on the server. It is understood that other details may also be listed regarding the files on the volume 118. It is further understood that the GUI 110 may provide a customizable interface in which an administrator can specify what types of summaries and histograms will be provided.

[0035] **Figure 5** illustrates a table listing information about directories on the server 102. The table 500 includes several columns, listing the directory name in the column 502, the number of files in the directory in the column 504, the total size of the files in the directory in the column 506, and the average time of the last access to files in the directory in the column 508. The agent 112 collects this information during the file walk, and compiles the table. The MMA 104, in many instances, does not have the resources to generate these tables or collect these results. This is especially true where there are several agents scanning a single storage server. Having the agents perform these tasks will save resources that the MMA 104 may require for other tasks.

[0036] The collected information about the directories on a storage server can be useful for several reasons. The administrator can find bottlenecks in the system, as well

as directories that have an abnormally large number of files or total size. In other embodiments, another table, similar to the table 500 may be generated. This table may include cumulative statistics that list the total number of files in a directory, including the total statistics for all embedded directories found within that directory. For example, the column 506 may list the total size of all files in a directory and in the directory's subdirectories.

[0037] The column 508 lists the average last access time for the files located in the listed directory. The column 508 lists a time stamp, in other words, an average time during which all files in the directory were last accessed. For example, if a directory contained five files, one most recently accessed today, one yesterday, one two days ago, one three days ago, and the last four days ago, the average access time would be sometime two days ago. This is useful so that an administrator can easily determine how active the particular directory is, and whether there are a large number of files that are not being regularly accessed. For example, it appears that there are a number of stale files in the directory '/u/users/a/Aaron/' since the average access time is over eighteen months ago.

[0038] **Figure 6** illustrates a histogram showing server usage of certain users. The histogram 600 demonstrates how much space each user is occupying on the volume 118. An administrator can use this data to determine whether one user is occupying an abnormally large amount of space. In one embodiment, the MMA 104 can use this information to revoke the user's ability to store any more files. For example, the users 'Aaron' and 'Gibson' are using much more storage space than the other users. The

administrator can target these users to increase the amount of free space on the server, if needed.

[0039] The histogram 600 may be personalized by the administrator. For example, in a system with many users, it may be difficult for the administrator to parse the histogram 600. Therefore, the histogram 600 may list the users with the highest usage first, or only those users that are using more than a specified amount of storage space. A histogram showing the usage of many users may allow an administrator to determine the approximate percentage of users that are using an abnormally large amount of server space. It is understood that the data represented in the histogram 600 may also be displayed in other forms, such as in table form.

[0040] **Figure 7** illustrates a histogram showing the types of files stored on a server. The histogram 700 can be useful to determine the typical usage of the server, and to point out improper usage. The histogram indicates several different types of files, including core files, executable files, text files, video files, audio files, photos, and database files. In another embodiment, the types of files may be listed by a file extension or other file identifier. For example, the histogram 700 may include a category for those files having an '.mp3' extension if an administrator wants to determine the amount of system space used by these files. It is understood that many different types of files may be reported in the histogram 700. It is further understood that the amount of usage may be represented as a percentage of total storage space.

[0041] As can be seen from the histogram 700, there are approximately 10GB of audio files and 6GB of video files stored on the server. Depending on the use of the server, this may or may not be a problem. For example, the server may be a web server

that hosts multimedia files. In this case, it would be appropriate to have this amount of media files compared to other files. However, if the server is a corporate server, it may be inappropriate for individual users to be stored video and audio files in their personal accounts. The histogram 700 can serve as an indication to the system administrator that action needs to be taken regarding these files. It is understood that an administrator may also specify that the file type is only displayed if the files of a specific type are occupying an abnormally large amount of space on the volume 118.

[0042] The techniques introduced above have been described in the context of a NAS environment. However, these techniques can also be applied in various other contexts. For example, the techniques introduced above can be applied in a storage area network (SAN) environment. A SAN is a highly efficient network of interconnected, shared storage devices. One difference between NAS and SAN is that in a SAN, the storage server (which may be an appliance) provides a remote host with block-level access to stored data, whereas in a NAS configuration, the storage server provides clients with file-level access to stored data. Thus, the techniques introduced above are not limited to use in a file server or in a NAS environment.

[0043] This invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident to persons having the benefit of this disclosure that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention. The specification and drawings are accordingly to be regarded in an illustrative, rather than in a restrictive sense.